Observations from Blind Stair Traversal on Cassie

Kevin Green, Jonah Siekmann, Helei Duan, Yesh Godse, Jeremy Dao, Alan Fern and Jonathan Hurst
Collaborative Robotics and Intelligent Systems Institute, Oregon State University
{*greenkev, siekmanj, duanh, godsey, daoje, afern, jhurst*}@*oregonstate.edu*

I. MOTIVATION

When legged robots are commonplace and useful in the real world with little supervision, they will have capable perception systems to localize and observe the environment around them. However, perception systems have inaccuracies, latency and may fail to observe changes in the environment like when a person misses a step while climbing a flight of stairs. This is a strong motivation to utilize a control method that is robust to variations in terrain properties. The extreme demonstration of this type of robustness is blind controllers, which are controllers that have no sensing of the world around them other than physical proprioception.

In this work we analyze the strategies of a recurrent learned controller that successfully ascends and descends flights of stairs without the use of perception. The controller transfers from simulation to hardware and can scale various flights of stairs, navigate large single step downs and traverse terrain well beyond the training environments (video of controller on hardware). The resulting controller is particularly interesting to analyze because our reward definition does not use any expert reference trajectories which could bias the behavior. We see recognizable strategies employed by humans [1] and animals [2] emerge from a relatively objective reward definition.

II. METHODS

We employ a similar approach to previous work that generated a wide variety of gaits on Cassie [3]. We utilize an LSTM recurrent neural network as the control policy and train it using a Mujoco simulation through the PPO algorithm. The reward signal is quite objective in terms of the resulting gait because it does not contain a reward for matching a reference trajectory. It consists of rewards for simple phase based foot force and velocity, stable body orientation, matching commanded velocity, low torques and smooth actions. The most important difference between this and the previous work is the environments in which the training occurs. We place in front of the robot an ascending flight of steps, followed by a plateau, then a descending flight of steps. To promote generalization we randomize the distance to the start of the stairs, the dimensions of the steps, the number of steps, the length of the plateau and the slope of the ground.

III. RESULTS

To understand the highly successful strategies that emerge we compare the stairs policy to a control policy trained on only flat ground. The stair policy has greater swing foot clearance and a more aggressive leg retraction policy while walking over flat ground (see Fig. 1). Further, we see that the foot swing path is significantly altered when the robot is in the process of descending or ascending a flight of stairs. The policies have an internal memory state which we postulate encodes an estimate of if the robot is on flat ground, ascending stairs or descending stairs (maybe even the dimensions of the stairs). This theory is supported by the change we observe in the swing foot profile when the robot is ascending or descending stairs. We also observe that the stair policies have a larger cost of transport ($0.46 \pm .03$) compared to flat ground policies ($0.38 \pm .02$) at 1 m/s on flat ground in simulation.

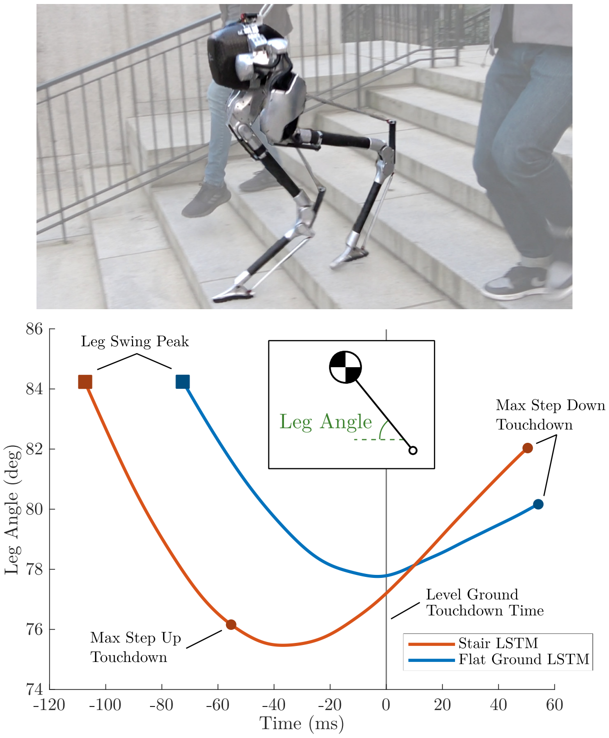

Fig. 1: (Top) Cassie descending a flight of stairs. (Bottom) A comparison of the swing leg angle of the Stair LSTM policy and the Flat Ground LSTM policy while locomoting at 1.0 m/s.

ACKNOWLEDGMENTS

This work was supported by DARPA contract W911NF-16-1-0002, NSF Grant No, IIS-1849343 and DGE-1314109. Thanks to Intel Labs for computing resources and support

REFERENCES

[1] Yvonne Blum, Susanne W Lipfert, Juergen Rummel, and André Seyfarth. Swing leg control in human running. *Bioinspiration & biomimetics*, 2010.

[2] Monica A. Daley and Andrew A. Biewener. Running over rough terrain reveals limb control for intrinsic stability. *Proceedings of the National Academy of Sciences*, 2006.

[3] Jonah Siekmann, Yesh Godse, Alan Fern, and Jonathan Hurst. Sim-to-Real Learning of All Common Bipedal Gaits via Periodic Reward Composition. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2021.